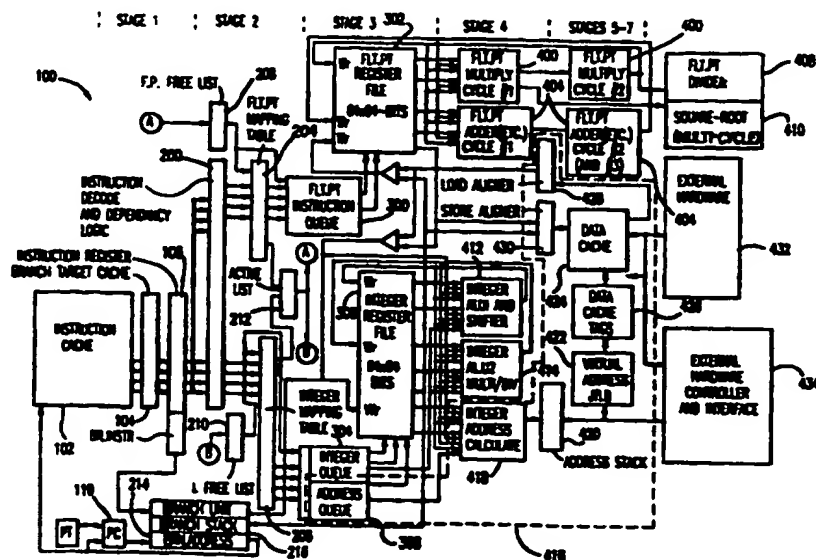




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 12/02		A1	(11) International Publication Number: WO 96/12228
			(43) International Publication Date: 25 April 1996 (25.04.96)
(21) International Application Number: PCT/US95/13300		(81) Designated States: JP, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 13 October 1995 (13.10.95)			
(30) Priority Data: 08/324,127 14 October 1994 (14.10.94) US		Published With international search report.	
(71) Applicant: SILICON GRAPHICS, INC. [US/US]; 2011 North Shoreline Boulevard, Mountain View, CA 94043-1389 (US).			
(72) Inventors: YEAGER, Kenneth, C.; Apartment 15-205, 1063 Morse, Sunnyvale, CA 94089 (US). KHURSHID, Mazin, S.; 3364 Bareoak Court, San Jose, CA 95148 (US).			
(74) Agents: KESSLER, Edward, J. et al.; Sterne, Kessler, Goldstein & Fox P.L.L.C., Suite 600, 1100 New York Avenue, N.W., Washington, DC 20005-3934 (US).			

(54) Title: REDUNDANT MAPPING TABLES



(57) Abstract

Redundant mapping tables (204', 206') for use in processors (100) that rename registers (302, 306) and perform branch prediction is presented. The redundant mapping tables (204', 206') include a plurality of redundant RAM cells (600). In the event of a branch instruction, the redundant RAM cells (600) can save the contents of the primary RAM cells (500) in a single clock cycle before the processor decodes and executes subsequent instructions along a predicted branch path. Should the branch instruction be mispredicted, the redundant cells (600) can restore the primary RAM cells (500) in a single clock cycle. A plurality of levels of redundant RAM cells (600) may be used to enable the nesting of a plurality of branch predictions at any one time.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

Redundant Mapping Tables

Background of the Invention

This invention relates in general to computers capable of performing register renaming and branch prediction and, in particular, to a computer capable of restoring register-renaming hardware to a pre-branch state in a single clock cycle in the event of a misprediction.

From the perspective of a programmer, instructions in a conventional processor are executed sequentially. When an instruction loads a new value into its destination register, that new value is immediately available for use by subsequent instructions. This is not true, however, for pipelined computer hardware because some results are not available for many clock cycles. Sequencing becomes more complicated in a superscalar processor, which has multiple execution pipelines running in parallel. But the hardware must behave as if each instruction were completed sequentially.

Each instruction depends on previous instructions which produced its operands, because it cannot begin execution until those operands become valid. These dependencies determine the order in which instructions can be executed. The actual execution order depends on the organization of the processor. In a typical pipelined processor, instructions are executed only in program order. The next sequential instruction may begin execution during the next cycle provided all its operands are valid. Otherwise, the pipeline stalls until the operands become valid. Because instructions execute in order, stalls usually delay all subsequent instructions. A sophisticated compiler can improve performance by re-arranging instructions to reduce the frequency of these stall cycles. In an in-order superscalar processor, several consecutive instructions may begin execution simultaneously, if all their operands are valid, but the processor stalls at any instruction whose operands are still busy. In an out-of-order superscalar processor, each instruction is eligible to begin execution as soon as its operands become valid, independently of the original instruction sequence. In effect, the

-2-

hardware re-arranges instructions to keep its execution units busy. This process is called "dynamic issuing."

Dynamic issue and execution of pipelined instructions creates a special need to monitor and resolve data dependencies between instructions. A newly-issued instruction is dependent on a previous instruction if, for example, the newly-issued instruction must use an output of the previous instruction as an operand. A technique called "register renaming" is used to resolve data dependencies which may affect the dynamic issue and execution of instructions. To ensure each instruction is given the correct operand values, logical register numbers (or "names") used in an instruction are mapped to physical registers maintained in hardware. Each time a new value is put in a logical register, it is assigned to a new physical register. Thus, although logical registers may have successive values simultaneously within a pipeline, a physical register has only a single value.

Although one or more instructions may begin execution during each cycle in a superscalar processor, each instruction takes several or many cycles to complete. Thus, when a branch instruction is decoded, its branch condition may not yet be known. However, through such well-known techniques as "branch prediction" and "speculative execution," a processor can predict whether the branch is taken, and then continue decoding and executing subsequent instructions along the predicted path. When this prediction is wrong, the processor must back up to the branch and take the other path. Whenever the processor discovers that it mispredicted a branch, it aborts all speculatively executed instructions to restore the processor to its state before the branch.

Restoration to a pre-branch state is a time consuming process, particularly when a number of instructions pending in the pipeline reference "renamed" registers. To restore these registers, typical superscalar pipeline processors sequentially "unname" each logical register in reverse order. This operation typically requires multiple clock cycles to complete. Such time consuming operation can significantly affect the overall performance of a processor since

branch instructions -- and therefore branch predictions -- are frequently encountered during the course of operation.

There is, therefore, a need for quickly returning a processor to a pre-branch state in the event of a branch misprediction when the processor employs register renaming.

Summary of the Invention

The present invention offers a highly efficient mechanism for saving and restoring register-renaming information to facilitate branch prediction and speculative execution. This mechanism enables the contents of register renaming mapping tables (i.e., memories used to rename logical register numbers to physical register numbers) to be saved and restored in a single clock cycle.

In one embodiment, the present invention provides a register-renaming mapping table having a plurality of primary RAM cells. A plurality of redundant RAM cells are coupled to these primary cells and capable of saving and restoring the contents of the primary RAM cells in a single clock cycle. Multiple levels of redundant RAM cells may be employed to save multiple versions of the contents of the primary RAM cells.

A better understanding of the nature and advantages of the present invention may be had with reference to the detailed description and the drawings below.

Brief Description of the Figures

Fig. 1 discloses a functional block diagram of a superscalar processor;

Fig. 2 illustrates a register mapping example;

Fig. 3 discloses a block diagram of primary and redundant mapping tables constructed according to the principles of the invention;

Fig. 4 discloses a single mapping table cell constructed according to the principles of the invention;

Fig. 5 discloses a primary bit cell constructed according to the principles of the invention;

5 Fig. 6 discloses a redundant bit cell constructed according to the principles of the invention;

Fig. 7 discloses a conventional RAM cell;

Fig. 8 discloses a block diagram of an integer free list; and

Fig. 9 is a timing chart showing register renaming timing.

10 ***Detailed Description of the Preferred Embodiment***

Contents

	I.	Superscalar Processor Architecture
	A.	Superscalar Processor Overview
	B.	Operation
15	II.	Register Operation
	A.	Register Renaming
	B.	Register Mapping Tables
	1.	Redundant Mapping Tables
	2.	Mapping Table Initialization
20	3.	Free Lists
	4.	Register Dependency
	5.	Register Map Timing
	6.	Active List
	III.	Branch Unit
25	I.	<i>Superscalar Processor Architecture</i>

Fig. 1 discloses a functional block diagram of a superscalar processor 100 which incorporates a redundant mapping mechanism built and operating in accordance with the present invention. As discussed below, this redundant

mapping mechanism enables efficient recovery from branch mispredictions. Processor 100, which generally represents the R10000 Super-Scalar Processor developed by Silicon Graphics, Inc., of Mountain View, California, provides only one example of an application for the redundant mapping mechanism of the present invention. This processor is described in J. Heinrich, *MIPS R10000 Microprocessor User's Manual*, MIPS Technologies, Inc. (1994).

A. *Superscalar Processor Overview*

A superscalar processor can fetch and execute more than one instruction in parallel. Processor 100 fetches and decodes four instructions per cycle. Each decoded instruction is appended to one of three instruction queues. These queues can issue one new instruction per cycle to each of five execution pipelines.

The block diagram of Fig. 1 is arranged to show the stages of an instruction pipeline and illustrates functional interconnectivity between various processor elements. Generally, instruction fetch and decode are carried out in stages 1 and 2; instructions are issued from various queues in stage 3; and instruction execution is performed in stages 4-7.

Referring to Fig. 1, a primary instruction cache 102 reads four consecutive instructions per cycle, beginning on any word boundary within a cache block. A branch target cache 104, instruction register 106 and instruction decode and dependency logic 200, convey portions of issued instructions to floating point mapping table 204 (32 word by 6 bit RAM) or integer mapping table 206 (33 word by 6 bit RAM). These tables carry out a "register renaming" operation, described in detail below, which renames logical registers identified in an instruction with a physical register location for holding values during instruction execution. A redundant mapping mechanism is built into these tables in accordance with the present invention to facilitate efficient recovery from branch mispredictions. Mapping tables 204 and 206 also receive input from a floating point free list 208 (32 word by 6 bit RAM) and an integer free list 210 (32 word

by 6 bit RAM), respectively. Output of both mapping tables is fed to active list 212 which, in turn, feeds the inputs of free lists 208 and 210.

A branch unit 214 also receives information from instruction register 106, as shown in Fig. 1. This unit processes no more than one branch per cycle. The branch unit includes a branch stack 216 which contains one entry for each conditional branch. As discussed below, processor 100 can execute a conditional branch speculatively by predicting the most likely path and decoding instructions along that path. The prediction is verified when the condition becomes known. If the correct path was taken, processing continues along that path. Otherwise, the decision must be reversed, all speculatively decoded instructions must be aborted, and the program counter and mapping hardware must be restored.

Referring again to Fig. 1, mapping tables 204 and 206 support three general pipelines, which incorporate five execution units. A floating-point pipeline is coupled to floating-point mapping table 204. The floating-point pipeline includes a sixteen-entry instruction queue 300 which communicates with a sixty-four-location floating point register file 302. Register file 302 and instruction queue 300 feed parallel multiply unit 400 and adder 404 (which performs, among other things, comparison operations to confirm floating-point branch predictions). Multiply unit 400 also provides input to a divide unit 408 and square root unit 410.

Second, an integer pipeline is coupled to integer mapping table 206. The integer pipeline includes a sixteen-entry integer instruction queue 304 which communicates with a sixty-four-location integer register file 306. Register file 306 and instruction queue 304 feed two arithmetic logic units ("ALU"); ALU#1 412 (which contains an ALU, shifter and integer branch comparator) and ALU#2 414 (which contains an ALU, integer multiplier and divider).

Third, a load/store pipeline (or load/store unit) 416 is coupled to integer mapping table 206. This pipeline includes a sixteen-entry address queue 308 which communicates with register file 306. The address queue 308 is preferably a superscalar processor having the capability to track memory-dependencies of

memory-access instructions. In accessing a two-way set-associative data cache, the address queue 308 imposes a dependency for accesses to the same cache set to prevent unnecessary cache thrashing.

5 More particularly, the address queue 308 is a highly efficient apparatus for tracking memory dependencies of memory-access instructions that may be executed out of order. This apparatus also provides for special identification of portions of a memory cache set to prevent unnecessary cache thrashing.

10 The address queue 308 holds a plurality of entries used to access a set-associative data cache. This queue 308 includes a comparator circuit, first matrix of RAM cells and second matrix of RAM cells. The comparator circuit compares a newly calculated partial address derived from a new queue entry with a previously calculated partial address derived from one of a number of previous entries. The first matrix of RAM cells tracks all of the previous entries in the queue that use a cache set that is also used by the new queue entry. The second
15 matrix of RAM cells tracks queue entries that are store instructions which store a portion of data in the data cache which is accessed by a subsequent load instruction. The address queue 308 may also assign status bits to certain blocks stored in the cache to identify the type of access allowed; i.e., random or sequential.

20 Register file 306 and address queue 308 feed integer address calculate unit 418 which, in turn, provides virtual-address entries for address stack 420. These virtual addresses are converted to physical addresses in joint translation lookaside buffer (JTLB) 422, and used to access a data cache 424 and data cache tags 426.

25 The JTLB 422 is a translation buffer for detecting and preventing conflicting virtual addresses from being stored therein. The JTLB 422 is a highly efficient mechanism for implementing translations of virtual memory addresses into physical memory. The JTLB 422 accurately prevents identical virtual page addresses from being stored in the translation buffer without having to shut down and reset the translation buffer. The JTLB 422 includes a first register for storing

a first physical page value, a second register for storing a second physical page value, a third register for storing the value representing the page size, a fourth register for storing a virtual address corresponding to the values stored in the first and second registers, and a comparator for comparing a current virtual address with the value in the third register. The comparator generates a signal if the compared values are equal. This signal causes either the first or second register to be read.

Data input to and output from data cache 424 pass through store aligner 430 and load aligner 428, respectively.

The data cache 424 is a memory array for storing instructions and data which is used by the CPU. Employing interleaved techniques, a high hit rate is achieved to effectively reduce the number of accesses to slower main memory. Multiplexing circuits enable the memory array to be more densely packed and implemented with lower numbers of sense amplifiers.

The cache 424 preferably operates in parallel with the translation lookaside buffer to reduce its latency. The cache 424 contains two 2-way set-associative arrays that are interleaved together. Each 2-way set-associative array includes two arrays, one each for the tag and data. By having four independently operating cache arrays, up to four instructions can operate simultaneously. The bits in each data array are interleaved to allow two distinct access patterns. For example, when the cache 424 is loaded or copied back, two double words in the same block are accessed simultaneously. When the cache 424 is read, the same doubleword location is simultaneously read from both blocks with the set. Further, by using a multiplexer, the number of sense amplifiers for reading and writing are reduced, thereby saving significantly valuable space on the die.

Address stack 420 and data cache 424 also communicate with external hardware controller and interface 434. Further, data cache 424 and controller/interface 434 communicate with external hardware 432.

B. Operation

Processor 100 uses multiple execution pipelines to overlap instruction execution in five functional units. As described above, these units include the two integer ALUs 412, 414, load/store unit 416, floating-point adder 404 and floating-point multiplier 400. Each associated pipeline includes stages for issuing instructions, reading register operands, executing instructions, and storing results. There are also three "iterative" units (i.e., ALU#2 414, floating-point divide unit 408 and floating-point square root unit 410) which compute more complex results.

Register files 302 and 306 must have multiple read and write ports to keep the functional units of processor 100 busy. Integer register file 306 has seven read and three write ports; floating-point register file 302 has five read and three write ports. The integer and floating-point pipelines each use two dedicated operand ports and one dedicated result port in the appropriate register file. Load/Store unit 416 uses two dedicated integer operand ports for address calculation. It must also load or store either integer or floating-point values, sharing a result port and a read port in both register files. These shared ports are also used to move data between the integer and floating-point register files.

In a pipeline, the execution of each instruction is divided into a sequence of simpler operations. Each operation is performed by a separate hardware section called a stage. Each stage passes its result to the next stage. Usually, each instruction requires only a single cycle in each stage, and each stage can begin a new instruction while previous instructions are being completed by later stages. Thus, a new instruction can often begin during every cycle.

Pipelines greatly improve the rate at which instructions can be executed. However, the efficient use of a pipeline requires that several instructions be executed in parallel. The result of each instruction is not available for several cycles after that instruction entered the pipeline. Thus, new instructions must not depend on the results of instructions which are still in the pipeline.

Processor 100 *fetches* and *decodes* instructions in their original program order, but may *execute* and *complete* these instructions out of order. Once completed, instructions are "*graduated*" in their original program order. Instruction fetching is carried out by reading instructions from instruction cache 102, shown in Fig. 1. Instruction decode operation includes dependency checks and register renaming (discussed below), performed by instruction decode and dependency logic 200 and mapping tables 204 or 206, respectively. The execution units identified above compute an arithmetic result from the operands of an instruction. Execution is complete when a result has been computed and stored in a temporary register identified by register file 302 or 306. Finally, graduation commits this temporary result as a new permanent value.

An instruction can graduate only after it and all previous instructions have been successfully completed. Until an instruction has graduated, it can be aborted, and all previous register and memory values can be restored to a precise state following any exception. This state is restored by "unnaming" the temporary physical registers assigned to subsequent instructions. Registers are unnamed by writing an old destination register into the associated mapping table and returning a new destination register to the free list. Renaming is done in reverse program order, in the event a logical register was used more than once. After renaming, register files 302 and 306 contain only the permanent values which were created by instructions prior to the exception. Once an instruction has graduated, however, all previous values are lost.

Active list 212 is a list of "active" instructions in program order. It records status, such as which instructions have been completed or have detected exceptions. Instructions are appended to its bottom when they are decoded. Completed instructions are removed from its top when they graduate.

II. Register Operation

Processor 100 uses "register renaming" to manage dependencies between registers. Renaming is implemented with mapping tables 204 and 206, free lists 208 and 210, floating-point register file 302, integer register file 306, and active list 212. The free lists contain the names of all physical registers (i.e., locations in register files 302 and 306) that are currently not in use. Active list 212 contains an entry for each instruction being processed in a pipeline. Instructions are appended to the active list as they are decoded and are removed as they graduate. As noted above, active list 212 lists all pending instructions in order.

A. Register Renaming

Register renaming is a technique used to keep track of changing register values. Processor 100 generates a myriad of temporary register results as it executes instructions. These temporary values, along with permanent values, are stored in register files 302 or 306 along with permanent values. Temporary values become new permanent values when the corresponding instructions graduate. Renaming identifies where the current (temporary or permanent) value of each register is in the register file.

Renaming distinguishes between logical registers, which are referenced within instruction fields, and physical registers, which are locations in the hardware register files. The programmer is aware of only the logical registers; the implementation of physical registers is entirely hidden. Only physical registers are actually implemented in hardware. Logical register numbers are dynamically mapped into physical register numbers. This mapping is implemented using mapping tables 204 and 206, as shown in Fig. 1, which are updated after each instruction is decoded. Each new result is written into a new physical register. Its value is temporary; the previous contents of each logical register can be restored, in case its instruction must be aborted following an

-12-

exception or a mispredicted branch. In the event of an exception, active list 212 is used to restore logical register contents. Alternatively, in the event of a mispredicted branch, redundant mapping tables as described below are used to restore the pre-branch state of logical registers.

5 Renaming simplifies dependency checks. Logical register numbers can be ambiguous when instructions are executed out of order because the same logical register may be repeatedly loaded with different values if it is specified by several instructions that are simultaneously in the pipeline. Because physical register numbers uniquely identify each result, dependency checking is
10 unambiguous. The queues and execution units use physical register numbers. In accordance with Fig. 1, integer and floating point registers are implemented with separate renaming maps 204, 206 and multi-port register files 302, 306.

 Instructions are executed "dynamically" after resolving all dependencies on previous instructions; i.e., each instruction can begin operation as soon as all
15 its operands are available. As such, an instruction need not wait until every previous instruction has been executed. A busy-bit table (not shown) keeps track of which physical registers are busy (i.e., not available). Each register is initially busy when it is assigned from the associated free list; it becomes not busy when a result is stored.

20 If the physical register number of an operand is not available when the instruction is decoded, that instruction must wait in a queue. The register number of the operand is compared to the destination register number of each instruction that is completed. When it matches, the operand becomes available. Because physical register numbers are uniquely assigned to each instruction, this match
25 can occur only for the instruction which generated the required result. With a few exceptions, integer queue 304 and address queue 308 have integer operand registers, and floating-point queue 300 has floating-point operand registers.

 Fig. 2 illustrates the mapping of a simple integer ALU instruction such as an ADD. ADD instruction 250 is configured in a preliminary format while being
30 held in a main memory. The instruction fields are rearranged prior to being

stored in instruction cache 102 through a "predecode" operation 251. Logical registers in processor 100 are usually selected using five-bit numbers 252 (rs), 254 (rt) and 256 (rd), shown in Fig. 2. The hardware stores copies of these registers in "physical registers" held in register files 302 and/or 306, shown in Fig. 1. Each physical register is identified by a six-bit physical number in the appropriate register file. The existence of these physical registers, and the mapping of logical registers to physical registers, are invisible to the programmer.

In addition to logical register numbers, instruction 250 contains function-related fields 258 (opcode), 260 (function), and 262 (subfunction).

Logical operand register numbers 252 and 254 are mapped directly to physical register numbers held in mapping table 206, resulting in physical operand register numbers 264 (OpA) and 266 (OpB). (Physical operand register number 268 (OpC) is used only in floating-point operations.) Conversely, logical destination register number 256 is mapped to a physical destination register number held in integer free list 210, as shown by line 270, resulting in a physical register number 272 (Dest).

In addition, function-related fields 258, 260 and 262 undergo a further predecode operation that groups these fields into a single function field 274. This function field and the foregoing physical numbers are gathered in a single mapped instruction 276, which is placed in integer queue 304 to await its turn for execution.

In addition to mapping logical destination register number 256, free list 210 forwards physical register number 272 to mapping table 206, as shown by line 278, thereby updating the table. As such, mapping table 206 reflects the latest mapping of logical destination register number 256. The old physical destination 282 associated with logical number 256 is output from mapping table 206 and appended to active list 212, as shown by line 280. Logical address 256 is also appended to active list 212, as shown by line 284. These two values, along with done bit 286 and exception field 288, form the latest entry in active list 212, which functions as a first-in-first-out ("FIFO") list. Old physical destination 282

will ultimately be returned to free list 210. as indicated by line 286. when instruction 276 graduates.

Clearly, there are more physical than logical registers since physical registers also store tentative results for instructions which have completed but have not yet graduated. A logical register may have a sequence of values as instructions flow through the pipeline. Whenever an instruction modifies a register, a new physical register is assigned to the logical destination register. These assignments are stored in register mapping tables 204 and 206. As each instruction is decoded, each of its logical register fields is replaced with the corresponding physical register number.

Each physical register is written exactly once. After it is output from a free list and until it is written with a result from a corresponding instruction, the physical register is "busy". If a subsequent instruction needs the values of this physical register, such instruction must wait until the register is written. After being written, the register is "ready", and its value is never changed. When a subsequent instruction changes the corresponding logical register, the result will be written into a new physical register. When this subsequent instruction graduates, the old value is no longer needed, and the old physical register becomes "free" for re-use. Thus, physical registers have unambiguous values.

Integer and floating-point registers are stored in separate register files and are renamed independently. Processor 100 contains 33 logical and 64 physical integer registers. (The extra logical register is the result of predefined register numbers associated with certain arithmetic operations.) Processor 100 also contains 32 logical and 64 physical floating-point registers.

B. Register Mapping Tables

As shown in Fig. 1, processor 100 includes two register mapping tables; integer mapping table 206 and floating-point mapping table 204. In this processor, each table maps a five-bit logical register number into a corresponding

six-bit physical register number. The integer table requires 33 six-bit entries, or 198 bits total. The floating-point table requires 32 six-bit entries, or 192 bits total.

5 Referring to Fig. 3, four write ports 350 are used to rename physical destination registers of the four instructions being simultaneously decoded (i.e., a total of four values being output from integer free list 208 and/or floating-point free list 210, as shown in Fig. 1). The same ports are used for both mapping tables 204 and 206, because an instruction may have either an integer or a floating result. Write ports 350 are also used by branch stack 216 to "update" mapping
10 tables 204 and/or 206 when these tables are restored from redundant mapping tables 204' and 206'. The need for this updating operation is described below. If a six-bit word is written by any write port during the update operation, it is not restored from a redundant mapping table.

15 Integer mapping table 206 has twelve read ports 354, as shown in Fig. 3, which map three logical register numbers (i.e., operands 368 and 370 and destination 372) for each integer instruction 366 (up to a maximum of four instructions in parallel). Each line 374 represents a five-bit logical register number (maximum of twelve lines for four instructions) and each line 376 represents a six-bit physical register number.

20 Floating-point mapping table 204 has sixteen read ports 352, as shown in Fig. 3, which map four logical register numbers (i.e., operands 358, 360 and 362, and destination 364) for each floating-point instruction 356 (up to a maximum of four instructions in parallel). Each line 378 represents a five-bit logical register number (maximum of sixteen lines for four instructions) and each line 380
25 represents a six-bit physical register number. Instruction register 382 provides a total of four instructions (i.e., floating-point and/or integer) to the foregoing read ports during each cycle. As discussed below, each mapping table includes a number of mapping table cells (primary and redundant) and read and write decoders.

1. Redundant Mapping Tables

Each mapping table includes a redundant map containing four entries. Whenever a branch instruction is executed tentatively, integer mapping table 206 and floating-point mapping table 204 are copied in parallel into a selected
5 redundant entry. This copy is retained in case the branch was mispredicted. When a mispredicted branch is undone, the original contents of the mapping tables are restored from the corresponding redundant entry. Referring to Fig. 3, floating-point redundant mapping table 204' contains four 192-bit entries (6 bits
10 each for 32 registers). Integer redundant mapping table 206' contains four 198-bit entries (6 bits each for 33 registers). Each bit entry is held by a conventional RAM cell, described below. Although shown as separate units in Fig. 3 for purposes of discussion, mapping table 204 (206) and redundant mapping table 204' (206') are embodied in a single mapping table in the preferred embodiment. Accordingly, a "single" mapping table cell 700, as discussed below and shown
15 in Fig. 4, includes a primary portion 500 (corresponding to tables 204 or 206) and redundant portion 600 (corresponding to tables 204' or 206').

Each redundant entry is a copy of a mapping table at the beginning of the decode cycle (i.e., stage 2) for the predicted branch instruction. Because processor 100 decodes up to four instructions in parallel, a redundant copy does
20 not correspond exactly to the state of the mapping table after the branch instruction. The table must be updated with changes made for preceding instructions that were decoded during the same cycle. Branch stack 216 retains this information and updates the appropriate mapping table(s), as shown in Fig. 3, when redundant mapping tables are restored. The contents of branch stack 216
25 is discussed below.

Fig. 4 illustrates single mapping table cell 700 of mapping tables 204, 204' and 206, 206'. As shown in this figure, cell 700 includes a primary bit cell 500 coupled to redundant bit cell 600 through interconnections 520, 524 and 526.

This circuit includes a number of externally coupled data lines 502-509 and 522 and control lines 510, 512, 514, 516, 518, 660, 662, 664, 666, 668, 670, 672, and 674, which are discussed in greater detail in connection with Figs. 5 and 6.

Referring to Fig. 5, primary bit cell 500 includes a primary RAM cell 501
5 coupled to a communication network facilitating writing and reading operations. This network includes data input lines 502, 504, 506 and 508 coupled to input-output ("I/O") port 580 of cell 501 through transistors 530, 532, 534 and 536, respectively. Similarly, data input lines 503, 505, 507 and 509 are coupled to I/O
10 port 582 of cell 501 through transistors 538, 540, 542 and 544. Lines 503, 505, 507 and 509 convey complemented values of signals conveyed on lines 502, 504, 506 and 508, respectively. Control signals 510, 512, 514 and 516 are coupled to the gates of transistors 530/538, 532/540, 534/542, and 536/544, respectively.

As further shown in Fig. 5, data input line 520 is also coupled to I/O ports
15 580 and 582 of cell 501. Specifically, line 520 is coupled to serially disposed inverters 546 and 548 which are, in turn, coupled to I/O port 580 through transistor 570. Additionally, line 520 is coupled to inverter 550 which is, in turn, coupled to I/O port 582 through transistor 572. Control line 518 is coupled to the gates of transistors 570 and 572.

I/O port 580 is also coupled to output lines 524 and 526 through inverters
20 552, 554, 556, 558, 560 and 562, as shown in Fig. 5. Similarly, I/O port 582 is coupled to output line 522 through inverters 564, 566 and 568.

As discussed above, processor 100 decodes up to four instructions per
25 cycle. Any one of these four instructions may access cell 501 to identify a destination register using data retrieved from a free list, as described above. Lines 502, 504, 506 and 508 each can convey a single bit of data that is a portion of a destination register number associated with an instruction currently undergoing mapping operations. At the same time, lines 503, 505, 507 and 509 convey complemented values of lines 502, 504, 506 and 508, respectively, to enable writing of cell 501. If four instructions are being mapped, all eight lines
30 will be active. Control signals 510, 512, 514 and 516 -- derived from logical

destination register numbers held in instructions being mapped -- determine which data line will write to cell 501.

Fig. 7 details the structure of primary RAM cell 501, which includes two conventional CMOS inverters 750 and 760. Inverter 750 is constructed from transistor 710 coupled in series to transistor 712. The source of transistor 710 is coupled to VCC while the source of transistor 712 is coupled to ground. The drain of transistor 710 is coupled to I/O port 582 and delay circuit 718 which is, in turn, coupled to the gates of transistors 714 and 716. The source of transistor 714 is coupled to VCC and the drain is coupled to transistor 716, the gates of transistors 710 and 712, and I/O port 580. The source of transistor 716 is coupled to ground.

As is well known in the art, a low (high) signal applied to I/O port 580 and high (low) signal applied to I/O port 582 will result in a steady state wherein I/O port 580 is maintained at a low (high) level and I/O port 582 is maintained at a high (low) level. Hence, data lines 502-509 convey noncomplemented and complemented signals to cell 501 when writing a new state, as shown in Fig. 5.

Control lines 510, 512, 514 and 516 determine which signal will be allowed to pass to cell 501. In terms of an entire mapping table, these control lines identify which cells, if any, will be altered by up to four instructions currently undergoing mapping operations. As noted above, these signals are derived from logical destination register numbers held in instructions being mapped.

In addition to being altered by instructions undergoing mapping operations, cell 501 may also be written with data from redundant bit cell 600 through data line 520, as shown in Figs. 5 and 6. Control line 518 enables the passage of data on line 520 to I/O ports 580 and 582 of cell 501. This line is controlled by branch unit 214, which maintains a record of currently active branch operations and controls mapping table restoration.

Cell 501 may output data to other portions of processor 100 over line 522, or to redundant bit cell 600 over lines 524 and 526, as shown in Fig. 5.

Fig. 6 shows redundant bit cell 600, which includes redundant RAM cells 601, 603, 605 and 607 coupled to data input lines 524 and 526 through a number of transistors. (The structure of redundant RAM cells 601, 603, 605 and 607 is the same as primary RAM cell 501.) Specifically, control line 668 is coupled to the gates of transistors 638 and 640. Transistor 638 couples data line 524 to I/O port 680 of redundant RAM cell 601 while transistor 640 couples data line 526 to I/O port 682 of the same redundant RAM cell. As shown in Fig. 6, an identical configuration exists between control lines 670, 672 and 674 and redundant RAM cells 603, 605 and 607, respectively. Control lines 668, 670, 672 and 674 are coupled to control logic 695, which generates necessary control signals to identify which redundant RAM cell is to be written.

Logic 695 is part of branch unit 214, which can identify those redundant cells currently in use through information contained in branch stack 216 (i.e., this logic identifies available redundant RAM cells). The branch unit creates a branch mask of 4 bits; each set bit (i.e., logic 1) indicating that a particular redundant RAM cell is in use and each cleared bit indicating that a particular redundant RAM cell is available. The branch stack, as discussed below, contains one entry for each branch that is tentative. Each entry contains the necessary information to confirm the initial branch decision and to reverse the branch if that decision was guessed incorrectly. The stack has four entries, so speculative branches can be nested four deep. Accordingly, logic 695 selects the next available redundant RAM cell (from the choice of four) based on information contained in branch stack 216.

I/O port 680 of redundant RAM cell 601 is also coupled to inverter 610 which is, in turn, coupled to transistor 630. Transistor 630 is controlled by control line 660 which is coupled directly to the non-inverted gate of transistor 630 and indirectly to the inverted gate of this transistor through inverter 620. The output of transistor 630 is coupled to data output line 520 through inverter 618. As shown in Fig. 6, an identical configuration exists between control lines 662, 664 and 666 and redundant RAM cells 603, 605 and 607, respectively. Control

lines 660, 662, 664 and 666 are coupled to control logic 696, which generates necessary control signals to identify which redundant RAM cell is to be read. Again, this logic is part of branch unit 214, which generates signals based on information contained in branch stack 216 (i.e., branch mask) and the result of a branch verification operation, described below.

Referring to Fig. 3, current mapping tables 204 and 206 are stored in redundant tables 204' and 206' during branch instruction decoding. More specifically, at the beginning of a decode cycle for a predicted branch instruction (i.e., during pipeline stage 2), branch control logic 695 (Fig. 6) enables one of four control lines 668, 670, 672 or 674, thereby selecting one of four redundant RAM cells 601, 603, 605 or 607 to store a primary RAM cell 501. The value of primary RAM cell 501 is conveyed to redundant bit cell 600 via lines 524 (noncomplemented value) and 526 (complemented value). This operation is carried out in parallel during a single clock cycle for 198 primary RAM cells in integer mapping table 206 and 192 primary RAM cells in floating-point mapping table 204. The presence of four redundant RAM cells 601, 603, 605 and 607 for each primary RAM cell 501 enables the nesting of up to four speculative branches at any one time.

In the event of a misprediction, one of the four redundant RAM cells is accessed to restore the mapping table. More specifically, ALU#1 412 (Fig. 1) performs a comparison operation to verify an integer branch. Similarly, floating-point adder 404 (Fig. 1) performs a comparison operation to verify a floating-point branch. If either execution unit determines a misprediction has occurred, a restore signal is forwarded to branch unit 214 which controls mapping table restoration.

More specifically, processor 100 identifies which branches will be verified one cycle before it performs the operation. Up to four branches can be verified in any single cycle. During branch verification, only one set of redundant RAM cells (i.e., all cells associated with a particular branch) may be read in a

-21-

single cycle. Accordingly, the redundant RAM cells associated with the oldest branch being verified are read first (i.e., during verification).

If the oldest branch is mispredicted, primary RAM cells in mapping tables 204 and 206 are restored with the redundant RAM cells values in 204' and 206', respectively. Any subsequent branches are discarded (i.e., corresponding entries in stack 216 are deleted) since these are based on the oldest branch. Alternatively, if the oldest branch is predicted correctly, the associated entry in branch stack 216 is deleted and the redundant RAM cell values are discarded. In the following clock cycle, the set of redundant RAM cells associated with the oldest branch that has mispredicted (identified through verification operation in the previous cycle) are read and used to restore tables 204 and 206. If no mispredictions are identified, associated entries in branch stack 216 are deleted during the branch verification cycle, and the operation is complete.

Referring again to Fig. 6, branch unit control logic 696 enables one of four control lines 660, 662, 664 or 666, thereby selecting one of four redundant RAM cells 601, 603, 605 or 607 to output a recovered value to primary RAM cell 501 via line 520. Referring to Fig. 5, control line 518 (also controlled by branch unit 214) enables the recovered value to pass through transistor 570 (noncomplemented) and transistor 572 (complemented) into I/O ports 580 and 582, respectively, of primary RAM cell 501. Like the storing operation described above, restoration operation is carried out in parallel during a single clock cycle for 198 primary RAM cells in integer mapping table 206 and 192 primary RAM cells in floating-point mapping table 204.

2. *Mapping Table Initialization*

Mapping tables 204, 206 and free lists 208, 210 are initialized when the chip is reset, as summarized in Table 1. Each physical register number exists exactly once, either in the mapping tables or in the free lists. As the processor executes instructions, these numbers are moved from one list to the other.

COMPONENT	INITIALIZATION
Floating Point Mapping Table 204	Logical register n is mapped to physical register n, for n= 0...31.
Floating Point Free List 208	Location n selects register (n+32), for n= 0...31
Integer Mapping Table 206	Logical register n is mapped to physical register n, for n= 1...31
A first special execution-unit dependent register	Mapped to physical integer register 0
A second special execution-unit dependent register	Mapped to physical integer register 32
Integer Free List 210	Location n selects register (n+32), for n= 1...31

Table 1: Mapping Tables and Free Lists Initialization

3. Free Lists

Free Lists 208 and 210 contain lists of the available integer and floating-point registers, respectively. A "free" register is an unused physical register which may be assigned as the destination register of a newly decoded instruction.

Each Free List contains 32 entries. Each entry contains a 6-bit physical register number. These entries are read and written as a circular FIFO with up to four reads and four writes per cycle. To achieve this throughput, the lists have four eight-entry RAMs disposed in parallel. Referring to Fig. 8, integer free list 210 includes four-to-one input multiplexers (MUXes) 802, 804, 806 and 808 coupled to eight-entry RAMS 810, 812, 814 and 816, respectively. These MUXes direct old destination register numbers from active list 212. The outputs of these RAMs are coupled to latches 818, 820, 822 and 824 which are, in turn, coupled to output MUXes 826, 828, 830 and 832. Read pointer 834 and write pointer 836 identify entries in RAMS 810, 812, 814 and 816 via column-selection lines 837-844 (which identify entries in two columns) and row selection lines

845-847 (which select one of two identified columns for RAMS 810, 812 and 814).

Input MUXes 802, 804, 806 and 808 are coupled to four pairs of 4X1 control MUXes, represented collectively by a single control MUX 848. Control MUX 848 input is a product of graduation mask input 850, ActIntDest 852 and write select table 856 as illustrated in Fig. 8 and described in Table 2. Graduation mask 850 identifies which instructions graduate and ActIntDest 852 identifies which instructions have an integer destination. Write mask 855, output from MUX 854, identifies which row(s) (i.e., RAMs 810, 812, 814 and/or 816) is to be written. Selection control of MUXes 848 and 854 is carried out by the two least significant bits of write pointer 836; i.e., IFreeWPtr[1:0]. Table 2 assumes the next empty entry is in row 0 (i.e., IFreeWPtr[1:0] = 00). The Write Select outputs of Table 2 are rotated for other pointer alignments. Mask values of "1" indicate that an entry should be written in the corresponding row.

5

10

15

Graduate Mask	Write Mask	Write Select
0000	0000	xxxx
0001	1000	3xxx
0010	1000	2xxx
0011	1100	23xx
0100	1000	1xxx
0101	1100	13xx
0110	1100	12xx
0111	1110	123x
1000	1000	0xxx
1001	1100	03xx
1010	1100	02xx
1011	1110	023x
1100	1100	01xx
1101	1110	013x
1110	1110	012x
1111	1111	0123

Table 2: Free List Write Decoding

20

25

Similarly, output MUXes 826, 828, 830 and 832 are coupled to four pairs of control MUXes, represented collectively by a single control MUX 858. Control MUX 858 input is a product of the two least significant bits of read pointer 864, (i.e., IFreeRPtr[1:0]) fed through incrementer 866, as shown in Fig. 8 and further described in Table 3. Selection control of MUX 858 is carried out by destination mask input 860 and read select table 862. Like Table 2, Table 3 assumes the next available entry is in row 0 (i.e., IFreeRPtr[1:0] = 00). The inputs of MUX 858 are rotated for other pointer alignments.

IntDest Mask	Read Select
0000	xxxx
0001	xxx0
0010	xx0x
0011	xx01
0100	x0xx
0101	x0x1
0110	x01x
0111	x012
1000	0xxx
1001	0xx1
1010	0x1x
1011	0x12
1100	01xx
1101	01x2
1110	012x
1111	0123

Table 3: Free List Read Decoding

Each RAM (i.e., 810, 812, 814 and 816) needs to have only one read and one write port, because any four consecutive entries will be spread between the four RAMs. Read pointer 834 and write pointer 836 are each 5-bit values. The high three bits are decoded to select one of eight column-selection lines 837-844. The low two bits select one entry within one of two columns via row selection lines 845, 846 and 847. This is the first entry to be read or written. The other three are in sequential order, as indicated by the numbers enclosed within each RAM entry 868 of Fig. 8. This includes any entries directly below the addressed entry and entries above in the next column. This entry selection is implemented

by multiplexing between a single column-selection line via 2X1 MUXes 870 and row-selection lines 845-847 at each entry, as shown in Fig. 8.

When an instruction graduates, the "old" physical register number associated with its destination register is written back into the appropriate free list. (Fig. 8 shows integer free list 210.) This releases the physical register, so that it can be used again. To write each number into the correct RAM, each write port includes a 4-input MUX (i.e., 802, 804, 806 or 808, as shown in Fig. 8) which can select from any of the four graduating instructions. The controls for these multiplexers depend on graduation mask 850, ActIntDest 852 and the low two bits of the write pointer. The mask, which (together with signal 852) indicates which of the four graduating instructions is releasing an integer register, is decoded into four 2-bit signal pairs. The location of the leftmost destination register is written into the leftmost pair, etc. As noted above, this decoder -- as represented by Table 2 -- assumes that low write pointer bits are zero. The decode for other values in the pointer is generated by rotating these four select signals, based on the low two bits of the write pointer. (Note that the same FIFO organization is used within the active list. The first graduating instruction may be anywhere within the four words. This circuit does not need to keep the free registers in any particular order.)

Write pointer 836 is incremented by the number of instructions which graduate during each cycle.

During each cycle, the next four free list entries are read using read pointer 834. If one or more of the instructions being mapped require an integer destination register, these entries are allocated in sequence to those instructions. The decode table -- as shown in Table 3 -- assumes that the low pointer bits are zero. For other values, these decodes are translated using four pairs of 4-to-1 multiplexers, represented collectively by MUX 858. The input to these multiplexers is a cyclic series of 2-bit values, beginning with the low two bits of the read pointer (00-01-10-11-00 etc.). Note, unlike the write decoder, this read decoder controls the selects, not the inputs, of the control multiplexers.

Read pointer 834 is incremented by the number of free registers assigned during each cycle. These assignments may be speculative for instructions decoded after a branch instruction is predicted. When a branch is executed speculatively, the read pointer is written into its branch stack entry, described below. This pointer value is incremented by the number of instructions decoded, which will not be aborted if the branch is reversed.

4. *Register Dependency*

Four instructions can be decoded and renamed during each cycle. Each logical operand register in an instruction is mapped into physical register names using multiported mapping tables 204 and 206, as shown in Figs. 1 and 3. A logical destination register is renamed with a new physical name from free list 210 (if an integer) or 208 (if a floating point). However, if the logical destination register of any instruction is used as an operand of a subsequent instruction (within these four), the new physical name of that register is used instead of the name read from the map table.

The first of the four instructions has no operand dependencies. The second instruction may depend on the first. The third instruction may depend on either of the first two. The fourth instruction may depend on any of the first three. If an operand depends on more than one previous instruction, the destination of the later instruction is used. Conventional comparator circuitry is used to determine such dependencies.

5. *Register Map Timing*

Mapping tables 204 and 206 are read during phase 2 of the decode cycle, as shown in Fig. 9. These tables are updated by branch stack 216 with new register mapping during phase 1 of the following cycle.

-28-

When a speculative branch is decoded, mapping tables 204 and 206 are copied into redundant mapping tables 204' and 206' (Fig. 3), respectively, during phase 2 of the decode cycle. If that branch was mispredicted, the mapping tables may be restored by copying the redundant mapping tables during phase 2 of a branch restore cycle, occurring during an execution stage of processor 100.

6. *Active List*

Active list 212 contains an ordered list of all active instructions. An instruction becomes active after it has been decoded. It remains active until it graduates (completes in program order) or is aborted by reversed branch decision or an exception.

When an instruction with an exception graduates, all subsequent instructions are flushed from the pipeline. Active list 212 is emptied in reverse order. For each instruction, the logical register is renamed to its previous physical register, and its new physical register is returned to the free list.

Active list 212 is implemented as a 32-word circular FIFO, whose structure is similar to free lists 208 and 210. Like the free lists, active list 212 consists of four eight-entry RAM memories. Four consecutive entries can be written (during decode) or read (during graduation) by accessing one entry in each memory. Write data is aligned to the proper memory using four pairs of four-to-one multiplexers. When reading, however, most of the data is not aligned at the active list. New free registers are appended to free lists 208 or 210, and the entries are aligned using the input MUXes of the free lists (i.e., MUXes 802, 804, 806 and 808 of integer free list 210, for example), as shown in Fig. 8.

Done bit 286 and exception bit 288, as shown in Fig. 2, are initialized, usually to zero, when an instruction is decoded. These bits are set when an instruction completes execution. The "done" and "exception" bits control when instructions graduate. When "done," up to four instructions can graduate per cycle, provided no previous instruction had an exception.

-29-

Instructions graduate in original program order from active list 212, after they have been completed by a functional unit. Instructions complete (i.e., become eligible for graduation) when their results have been computed and stored in a register. Because execution and completion can occur out of order, each completion is recorded by setting done bit 286 for that instruction. Address queue 308 generates "done" signals for load and store instructions. For all other instructions, an execution unit sets the done bit of an instruction in active list 212. As instructions graduate, a "graduation unit" (not shown) removes instructions from active list 212 and appends their old destination registers to free lists 208 or 210 for re-use. Up to four instructions can graduate in parallel during each cycle.

If a functional unit detects an exception, it stores an exception code for the effected instruction in active list 212 rather than setting the done bit 286 of this instruction. This instruction is aborted and will not graduate, but there is no direct effect on other instructions in the pipeline. Instruction fetch and decode is suspended.

If an exception occurs on an instruction fetched after a mispredicted branch, the instruction will remain pending until that branch is reversed. Then, that instruction is aborted and the exception is ignored.

An exception is taken only when the corresponding instruction is ready to graduate, after all previous instructions have graduated. At this point, all subsequent instructions are aborted and the mapping tables 204 and 206 are restored to their state before the excepting instruction was decoded.

On a cycle after an exception is detected, the mapping tables are sequentially restored using the logical destination field 290 and old destination field 292 in active list 212, as shown in Fig. 2. Up to four registers are restored per cycle, beginning with the latest and working backward. Because there may be up to 32 instructions in active list 212, this restoration process may require up to eight cycles.

The "new" registers, which were assigned to the now aborted instructions when they were decoded, must be returned to the appropriate free list. This is

done by simply adjusting the read pointers of the two free lists, 208 and 210. It is not necessary to write into these lists because their original contents have not been overwritten. Because all instructions in the pipeline have been purged, all registers which are not in the mapping tables must be in the corresponding free list. For floating-point free list 208, there are thirty-two free registers (making the free list full). This is accomplished by resetting the read pointer to equal the write pointer. For integer free list 210, there are thirty-one free registers. The read pointer is set equal to the write pointer plus one.

III. Branch Unit

Processor 100 typically addresses instructions sequentially through the use of program counter (PC) 110, shown in Fig. 1. PC is typically incremented by one word (i.e., 4 bytes) after each instruction is executed. This sequential order is modified by jump and branch instructions, which load non-sequential values into PC 110. Jump instructions are unconditional changes in sequence. Branch instructions are conditional upon a selected test. For simplicity, the term "branch" as used in this text refers to either conditional or unconditional instructions.

Referring to Fig. 1, branch unit 214 decodes and initiates branches during pipeline stage 2. The value held in PC 110 is valid for stage 2 only; it points to the first instruction being decoded. The next value of PC is computed during stage 2.

Conditional branches create delays because the condition often cannot be tested for several (or many) cycles after the branch instruction is decoded. To reduce these delays, processor 100 predicts whether the branch will be taken and proceeds to execute instructions speculatively along the predicted path. If the prediction is correct, testing the condition has little effect on performance. However, when a branch is mispredicted, processor 100 must reverse direction and abort all speculatively decoded instructions. This processor must also restore

-31-

the pre-branch values held by PC 110 and register mapping tables 204 and 206, shown in Fig. 1. This is completed using branch stack 216 and redundant mapping tables 204' and 206', shown in Figs. 1 and 3, respectively. Processor 100 can handle up to four nested speculative branches and verify them in any order.

A branch instruction may be in any of the four words in instruction register (IR) 106, shown in Fig. 1. The earliest branch instruction is selected and then decoded. Only this one branch could be decoded during this cycle. If IR 106 contains a second branch instruction, decoding it and any subsequent instructions are delayed until the next cycle.

When a branch is reversed, active list 212 and mapping tables 204 and 206, shown in Fig. 1, must be restored to the point at which the execution paths diverged. That is, each unit is restored to its state after the last instruction which is executed in both the "branch taken" and "branch not taken" cases; defined as the "restore point."

As discussed above, redundant mapping tables 204' and 206' (Figs. 3-7) make a copy of mapping tables 204 and 206, respectively, for use if a branch is restored. Although processor 100 decodes up to four instructions in parallel, mapping tables 204 and 206, in their preferred embodiment, are not updated for each instruction individually. Rather, these tables make a copy at the beginning of the decode cycle (i.e., stage 2). Thus, the precise restore point for a particular branch instruction might not exist in redundant mapping tables 204' and 206'. As such, the changes made during the cycle are recorded separately in branch stack 216. That is, the destination registers of the four instructions decoded during the same cycle are saved in the branch stack.

If all four instructions are decoded in a single cycle, an alternative embodiment of redundant mapping tables 204' and 206' could make a new copy of mapping tables 204 and 206, respectively, during the following cycle. This embodiment would eliminate the need for saving the fourth destination register. However, this extra redundant mapping operation might coincide with a

-32-

redundant mapping operation for a second branch instruction. To accommodate simultaneous redundant mapping operations, stronger drivers would be required to accommodate simultaneous writes to redundant RAM cells.

A. Speculative Branches

5 As noted above, conditional branches can cause significant processing delays in a pipelined processor because the tested condition is often not determined for several cycles after the branch is decoded.

10 In processor 100, if the direction of a branch is unknown when the branch is decoded, the direction will be predicted and instructions will continue to be decoded speculatively. Instructions are fetched only along the predicted path, but up to four predictions can be nested. A branch is recorded in branch stack 216 until the prediction can be verified. Verification may be performed out of order.

B. Branch Prediction

15 The prediction of conditional branches depends on a branch history circuit contained in branch unit 214. This is an n-word by two-bit RAM that tracks how often each particular branch was taken in the past. The word is selected by certain bits of the address of the branch instruction. Each entry has a two-bit code, as indicated in Table 4. The two-bit code is updated whenever a final branch decision is made. The first bit is used to predict a branch taken (1) or not taken (0). The second bit helps decide when to change the prediction. The
20 prediction is changed only if two contrary branch decisions occur in a row.

State	Next State		Prediction
	Taken	Not	
00	01	00	Not taken. Last two branches were not taken.
01	10	00	Not Taken. Previous decision was contrary ("taken").
10	11	01	Taken. Previous decision was contrary ("not taken").
11	11	10	Taken. Last two branches were taken.

Table 4: Branch history RAM

C. Branch Stack

Branch stack 216, identified in Fig. 1, contains an entry for each branch instruction which is being executed speculatively. It contains four entries, so up to four speculative branch decisions can be nested. An entry is created whenever a branch path is predicted. The entry is deleted when the branch decision has been verified. Entries are also deleted if any previous branch was predicted incorrectly.

Each branch stack "entry" contains the information required to restore the state of processor 100 if a speculative branch is reversed. An entry consists of several physically distinct sections (i.e., held in various locations within the processor), which are associated with various data structure whose value must be restored. Several sections of a branch stack entry are identified in Table 5.

Section	Signal	Description
Branch Stack	BSBrMask	Branch mask of previous speculative branches.
Redundant Mapping Table	ShadowN	Copy of integer mapping table 206 (33 X 6 bits).
	ShadowF	Copy of floating-point mapping table 204 (32 X 6 bits).
	ShadowUp	Shadow map updates -- update four newly decoded registers in both mapping tables.
Active List	BSActWPtr	Active list - write pointer.
Free Lists	BSFreeNRPtr	Integer free list - read pointer.
	BSFreeFRPtr	Floating-point free list - read pointer.
PC	BSAltPC	Alternate branch path address.

Table 5: Branch stack entry

The following is a brief description of the entries and terms used in Table

5.

A "branch mask" indicates if each entry is dependent on any previous speculative branch. This indicates the original program order of the branches in branch stack 216.

Redundant mapping tables 206' and 204' make a copy of integer mapping table 206 and floating-point mapping table 204 during phase 2 of stage 2, respectively (see Fig. 9). As discussed above, these redundant tables are built within primary mapping tables (a map RAM structure) so that the entire contents can be copied in parallel in a single cycle. These mapping tables are also referred to as "shadow maps."

"Shadow map updates" retain the necessary information to update redundant mapping tables to reflect any changes made for instructions decoded during the cycle when the subject branch was originally decoded. These changes are effected when a redundant mapping table is restored. Each such instruction may change one physical (integer or floating-point) register. Each register is

specified with its six-bit physical number and a two-bit validity code. The code is "0" if the instruction has no destination register or if it was decoded after the "restore point." Codes "2" and "3" indicate the restored redundant mapping table must be updated for integer and floating-point destination registers, respectively.

5 Write pointer for active list 212 is used to append newly decoded instructions to the active list. When a speculative branch is reversed, all instructions after that branch are aborted. They are removed from active list 212 by restoring its write pointer.

10 Read pointers for free lists 208 and 210 are used to read the physical register numbers of the next available integer or floating-point registers. When a speculative branch is reversed, these registers are returned to the free lists by restoring their read pointers. The contents of these locations in the free lists has not been altered.

15 Alternate branch path address contains the next instruction address along the execution path not taken. It is reloaded into PC 110 (Fig. 1) if the branch prediction was wrong. If the branch is initially taken, the alternate address is the instruction after the "restore point," defined above. Otherwise, the alternate address is the target address of the unused branch instruction.

20 While the above is a complete description of the preferred embodiment of the invention, various modifications, alternatives and equivalents may be used. Therefore, the above description should not be taken as limiting the scope of the invention which is defined by the appended claims.

What Is Claimed Is:

- 1 1. In a processor capable of executing instructions and renaming
2 logical register numbers contained in said instructions to physical register
3 numbers, a register-renaming mapping table comprising:
4 a plurality of primary RAM cells; and
5 a plurality of first redundant RAM cells coupled to said plurality of
6 primary RAM cells through data output lines and data input lines, said data
7 output lines moving first data in parallel from said plurality of primary RAM cells
8 to said plurality of first redundant RAM cells in a first clock cycle, and said data
9 input lines capable of moving said first data in parallel from said plurality of first
10 redundant RAM cells to said plurality of primary RAM cells in a second clock
11 cycle.
- 1 2. The register-renaming mapping table of claim 1 further comprising
2 a plurality of second redundant RAM cells coupled to said plurality of primary
3 RAM cells through said data output lines and said data input lines, said data
4 output lines capable of moving second data in parallel from said plurality of
5 primary RAM cells to said plurality of second redundant RAM cells in a third
6 clock cycle, and said data input lines capable of moving said second data in
7 parallel from said plurality of second redundant RAM cells to said plurality of
8 primary RAM cells in a fourth clock cycle.
- 1 3. The register-renaming mapping table of claim 2 further comprising
2 a plurality of third and fourth redundant RAM cells coupled to said plurality of
3 primary RAM cells through said data output lines and said data input lines.

-37-

1 4. A pipeline processor comprising:
2 means for issuing instructions containing floating-point logical register
3 numbers and integer logical register numbers;
4 a first renaming means, coupled to said issuing means, for renaming
5 floating-point logical register numbers to physical register numbers;
6 a first storing means, coupled to said first renaming means, for storing
7 contents of said first renaming means in a first clock cycle in response to a branch
8 instruction;
9 a second renaming means, coupled to said issuing means, for renaming
10 integer logical register numbers to physical register numbers;
11 a second storing means, coupled to said second renaming means, for
12 storing contents of said second renaming means in said first clock cycle in
13 response to said branch instruction;
14 a branch prediction means, coupled to said first and second renaming
15 means, for predicting said branch instruction;
16 a comparator means, coupled to said branch prediction means, for
17 determining whether branch prediction is correct; and
18 a branch control means, coupled to said branch prediction means, which
19 enables said first and second storing means to restore said contents of said first
20 and second renaming means, respectively, in a single clock cycle when a branch
21 misprediction occurs.

1 5. The pipeline processor of claim 4 further comprising a free list
2 means, coupled to said first renaming means, for providing new mapping
3 information to said first renaming means.

1 6. The pipeline processor of claim 4 further comprising an active list
2 means, coupled to said first renaming means, for receiving old mapping
3 information from said first renaming means.

-38-

1 7. The pipeline processor of claim 5 further comprising a stack
2 means for recording and restoring said new mapping information provided by
3 said free list means, wherein said first storing means stores contents of said first
4 renaming means early in said first clock cycle, and said free list means provides
5 said new mapping information to said first renaming means later in said first
6 clock cycle.

1 8. The pipeline processor of claim 4, wherein said first storing means
2 includes means for holding a plurality of versions of said contents of said first
3 renaming means at any one time.

1 9. In a pipeline processor capable of renaming registers and
2 speculatively executing instructions, a method for handling branch mispredictions
3 comprising the steps of:

4 storing contents of a register renaming mapping table during a first clock
5 cycle;

6 predicting a branch instruction;

7 executing instructions pursuant to said prediction;

8 verifying branch prediction; and

9 restoring said contents of said register-renaming mapping table in a
10 second clock cycle when a branch misprediction occurs.

1 10. The method of claim 9 further comprising the step of restoring
2 changes to said register renaming mapping table occurring during said first clock
3 cycle by instructions that are decoded during said first clock cycle and executed
4 whether or not branch is taken.

1/10

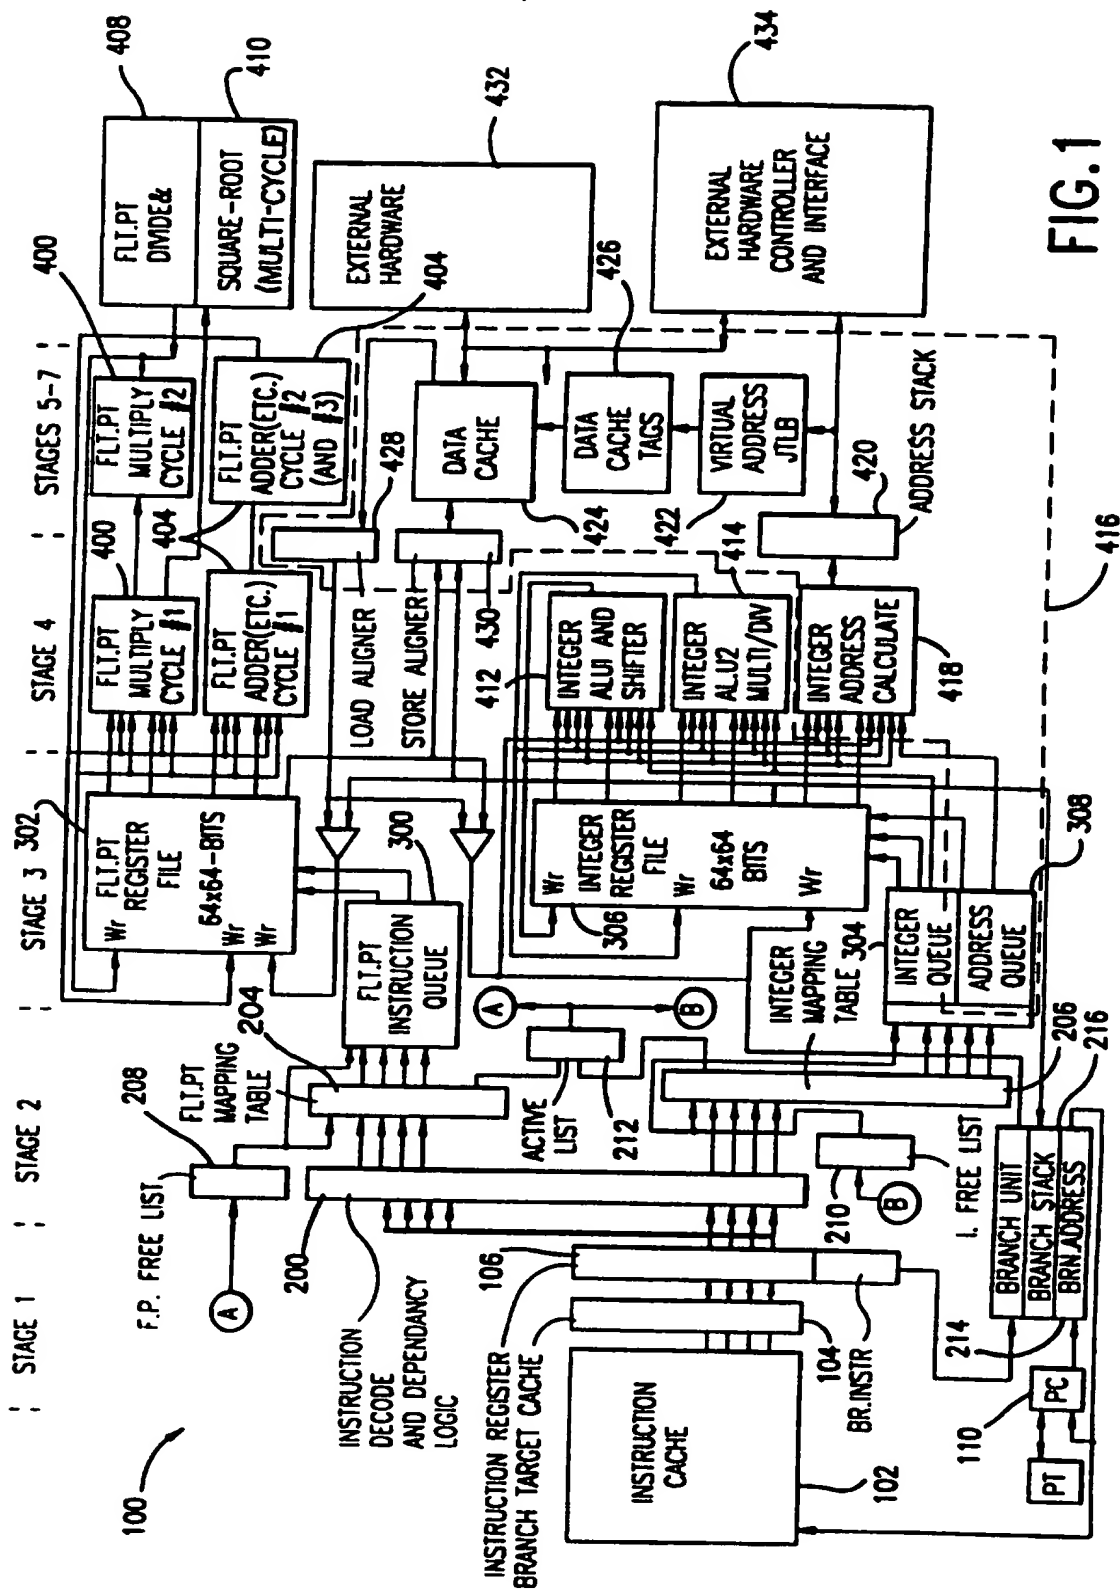



FIG. 1

2 / 10



SUBSTITUTE SHEET (RULE 26)

FIG.2




FIG.3




FIG. 4

5/10




FIG. 5

6/10

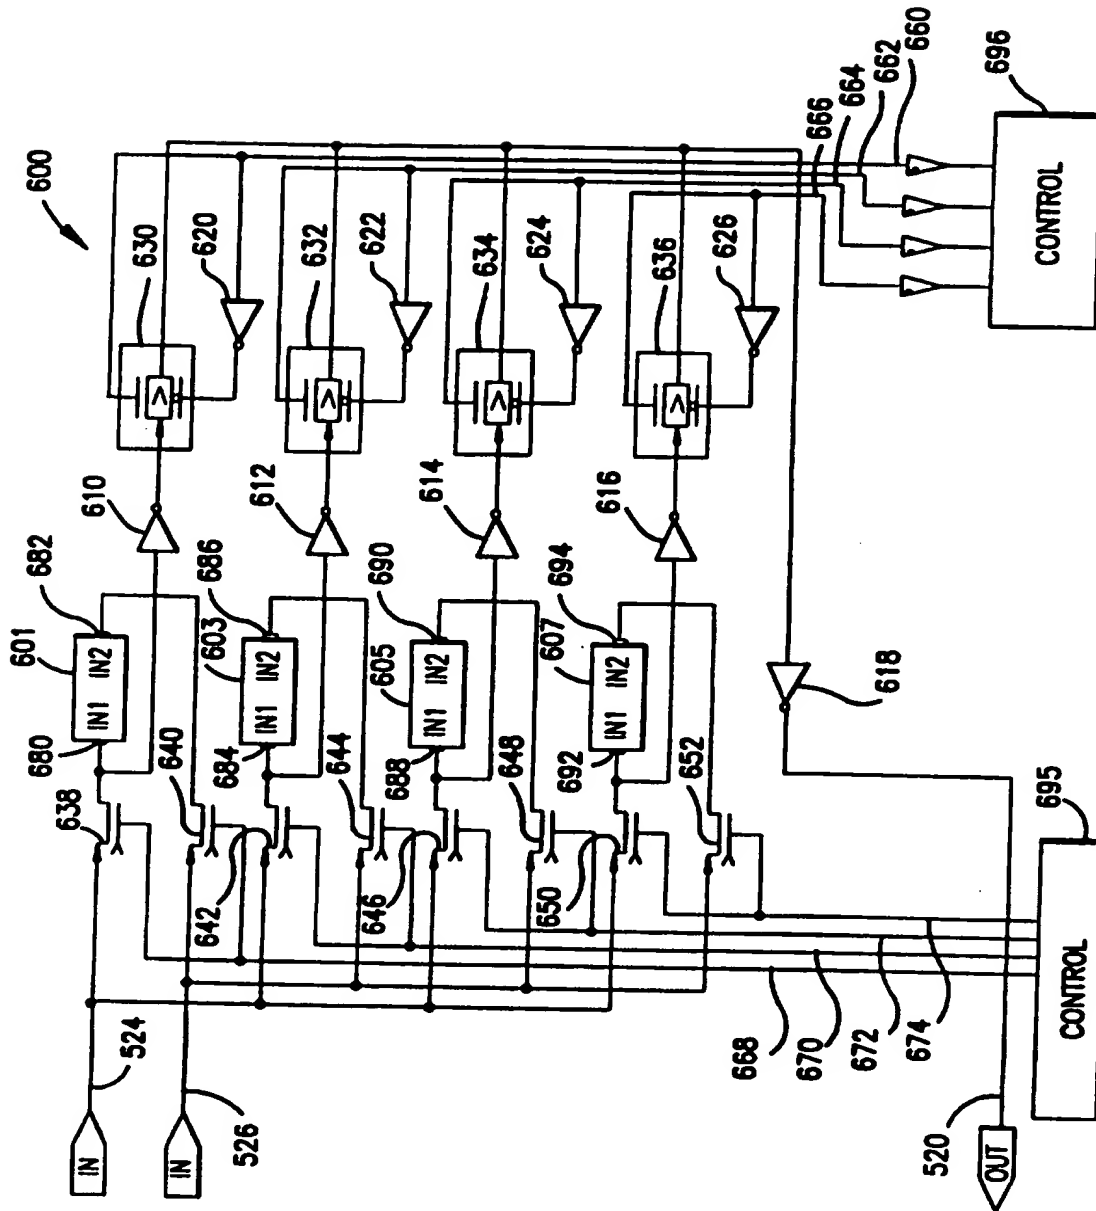


FIG. 6

7/10




FIG.7


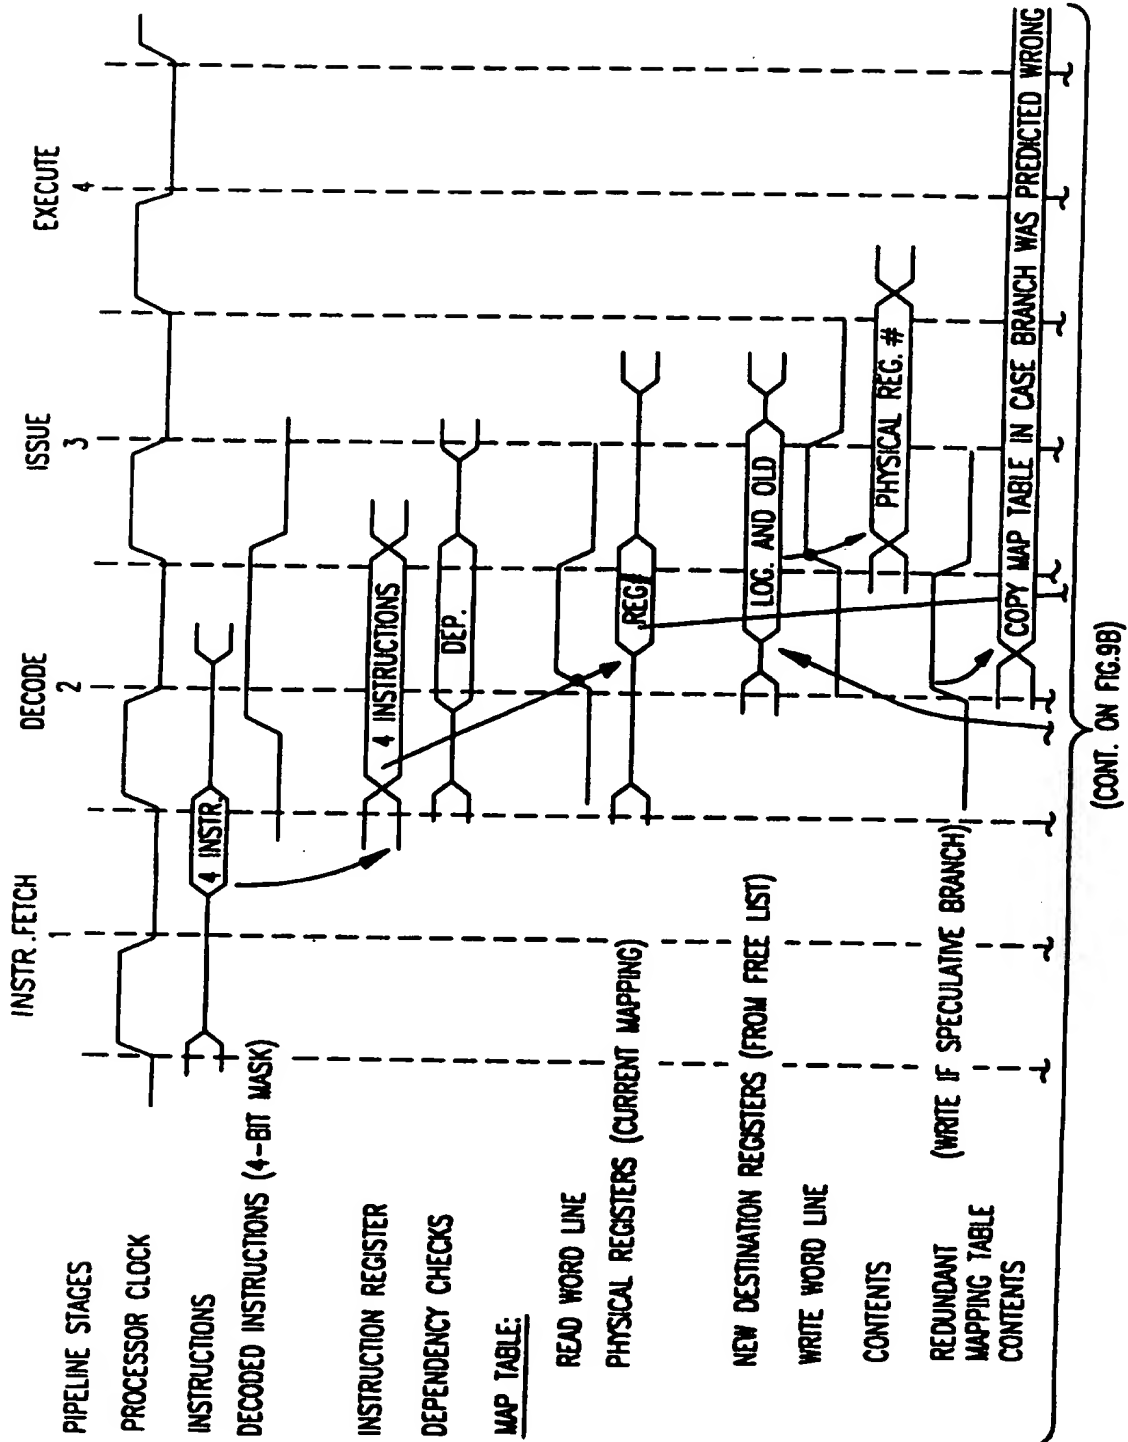


FIG. 8

9/10



(CONT. ON FIG.9B)

FIG.9A

10/10




FIG. 9B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US95/13300

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 12/02

US CL : 395/375

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 395/375, 401, 405, 414, 415, 416, 417, 418, 419, 420, 421.03, 421.06, 421.11, 427, 428, 488, 489, 496, 800

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
NONE

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS (U.S. Patents); APS (JPOABS)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US, A, 5,134,561 (LIPTAY) 28 JULY 1992, see abstract, figures 1A-5,, col. 7 (line 66 - et seq.), col. 8 (line 6 - et seq.)	1-10

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* "A"	Special categories of cited documents: document defining the general state of the art which is not considered to be part of particular relevance	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document published on or after the international filing date	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubt on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combinations being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"A"	document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

02 JANUARY 1996

Date of mailing of the international search report

01 FEB 1996

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

ROBERT B. HARRELL

Telephone No. (703) 305-9600

Jon Hill